

A Brief Review on the Development of HYPRE for GPUs

FASTMath4 Institute - All-Hands Meeting – Argonne National Laboratory



June 11

Ruipeng Li



LLNL-PRES-777450

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC




GPU support in HYPRE

- GPU support has been available in recent releases of HYPRE
 - enabled by various approaches including CUDA, OpenMP 4.5, RAJA, Kokkos
- Structured multigrid solvers: SMG, PFMG
 - Structured data types: grid, box, stencil, struct matrix and vector...
 - Computations are performed in BoxLoops
 - Completely ported to GPUs, MG setup and solve, in GPU device memory
- Unstructured algebraic multigrid: BoomerAMG
 - Unstructured data types: ParCSR matrices, Parvectors, ...
 - Solve phase has been ported: MATVEC, vector operations, and appropriate relaxations run on GPUs with unified memory
 - Setup phase: performed on CPUs with unified memory
 - Current focus: AMG setup phase on GPUs



Structured interface of HYPRE

- Provides access to the structured solvers: SMG, PFMG, ...
- Struct grid is composed of **boxes**



- Struct matrix and vector sit on top of struct grid
- Struct computations are performed via **BoxLoops**

HYPRE loop abstraction for computations: BoxLoops

- Example: $y := \alpha y$, y is a struct vector

```
1  hypre_ForBoxI(i, boxes) {  
2    ...  
3    hypre_BoxLoop1Begin(dim, loop_size, data_box, start, unit_stride,  
4      yi);  
5    [#pragma omp parallel for private(yi)]  
6    hypre_BoxLoop1For(yi) {  
7      yp[yi] *= alpha;  
8    }  
9    hypre_BoxLoop1End(yi);  
}
```

Different ideal memory access patterns

CPU: coarse-grained parallelism

1	1	1	1	1	1	1	1
1	1	2	2	2	2	2	2
2	2	2	2	2	3	3	3
3	3	3	3	3	3	3	3
		1	1	1	1	1	
		2	2	2	2	2	
		3	3	3	3	3	

CPU parallel BoxLoop

GPU: fine-grained parallelism

1	2	3	1	2	3	1	2
3	1	2	3	1	2	3	1
2	3	1	2	3	1	2	3
1	2	3	1	2	3	1	2
			1	2	3	1	2
			3	1	2	3	1
			2	3	1	2	3

GPU parallel BoxLoop

GPU BoxLoops have the same interface as the CPU ones

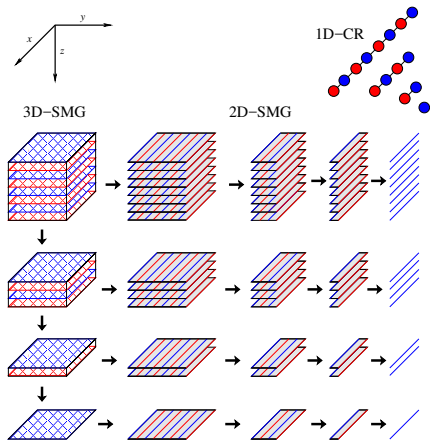
- Offload the OMP parallel region to GPU: with CUDA or OpenMP 4.5

```
1 #define hypre_BoxLoop1Begin(ndim,loop_size,dbox1,start1,stride1,i1) {\n2   /* host code: */ \n3   hypre_BoxLoopDeclareInit(ndim,loop_size) \n4   hypre_BoxKDeclareInit(1,start1,dbox1,stride1) \n5   /* device code (CUDA): */ \n6   BoxLoopforall(hypre_exec_policy,tot,HYPRE_LAMBDA(HYPRE_Int idx) {\n7       hypre_BoxLoopSet1(i1)
```

```
1   /* device code (OMP4.5): */ \n2   _Pragma(omp target teams distribute parallel for) \n3   for (thread=0; thread<tot; thread++) {\n4       hypre_BoxLoopSet1(i1)
```

- also available with RAJA or Kokkos



Complicated structured MG solvers can be ported seamlessly



- 3-D SMG semicoarsens in the z -direction and uses xy -plane smoothing by one V-cycle of 2-D SMG, which semicoarsens in the y -direction and uses 1-D line smoothing in the x -direction by Cyclic Reduction
- Interpolation operator is computed by performing a sequence of (simultaneous) SMG solves (of one dimension lower)
- SMG is recursively constructed and is applied in the solve phase
- $\text{SMG}(3\text{D}) \rightarrow \text{SMG}(2\text{D}) \rightarrow \text{CR}(1\text{D})$

Performance of SMG/PFMG on GPUs

- 3-D Poisson problem. Problem size $200^3 \times N_p$
- Running on **ray** at LLNL. IBM Power8 + 4 NVIDIA Tesla P100 (Pascal) per node. CUDA 9.2, IBM XL compilers





The solve phase of BoomerAMG has been ported to GPUs

- GPU kernels for MATVEC and vector operations have been implemented with CUDA or OpenMP 4.5
 - with the option of using cuSPARSE and cuBLAS
- Several GPU-appropriate smoothers have been implemented
 - L_1 Jacobi and polynomial smoothers
- Setup phase is much more complicated. Currently, remains on CPUs with CUDA unified memory

Performance of AMG on GPUs

- 3-D Poisson problem. Problem size $200^3 \times N_p$
- Running on **ray** at LLNL. IBM Power8 + 4 NVIDIA Tesla P100 (Pascal) per node. CUDA 9.2, IBM XL compilers




Work in progress: AMG setup on GPUs

- Main ingredients of classical AMG setup
 - Compute SoC matrix \mathbf{S} . Relatively easy to implement
 - Compute coarsening \mathbf{C}/\mathbf{F} . PMIS algorithm has been implemented
 - Compute interpolation \mathbf{P} . Direct interpolation has been implemented
 - Compute Galerkin product \mathbf{RAP} . Most difficult. Two algorithms have been implemented. Computed in pairs: $\mathbf{Q} = \mathbf{AP}$ and $\mathbf{RAP} = \mathbf{RQ}$

Distributed SpGEMM in ParCSR

- In parallel CSR, each process owns a slice of rows partitioned into diag and offd parts
- Decompose ParCSR $A \times \text{ParCSR } B$ into *local* CSR matrix operations: multiplication, (partial) addition, splitting, merging, and transposition for $A^T B$
- Also wrote CUDA kernels other than multiplication: with **Thrust**; much simpler
- (GPU) Communications are involved for sending/receiving *external* rows, and can be overlapped with computations





A complete hash-table based SpGEMM algorithm

► Compute SpGEMM in 4 steps

- ❶ Row NNZ estimation: to allocate “reasonable-sized” hash tables for ❷;
- ❷ Symbolic analysis: to compute row counts (bounds) and row pointers ic , and to allocate “adequate-sized” hash tables for ❸;
- ❸ Numeric multiplication: to compute the column indices and values in jc and c ;
- ❹ Post-processing: to remove the gaps between rows computed from ❸.

GPU ParCSR SpGEMM: 3-D 27-pt Laplacians

- BoomerAMG with HMIS coarsening and ext+i interpolation
- Compute $P^T A P$ on the 1st level of AMG
- Weak scalability study. Local problem size: 128^3
- ray: IBM Power8 + 4 NVIDIA P100. lassen: IBM Power9 + 4 NVIDIA V100
- On ray 4 – 64 GPUs (left) and lassen 4 – 128 GPUs (right). y-axis: time





HYPRE's memory model

- Three conceptual memory locations
 - HYPRE_MEMORY_HOST, HYPRE_MEMORY_DEVICE, HYPRE_MEMORY_SHARED
- Mapped to different physical memory in different configurations

	HYPRE_MEMORY_HOST	HYPRE_MEMORY_DEVICE	HYPRE_MEMORY_SHARED
c1	HOST	HOST	HOST
c2	HOST	CUDA DEVICE	CUDA DEVICE
c3	HOST	CUDA DEVICE	CUDA MANAGED

- c1: non-GPU configuration
 - c2: `--with-cuda, --with-device-openmp`
 - c3: c2 + `--enable-unified-memory`
- `hypr_TAlloc(HYPRE_Complex, n, HYPRE_MEMORY_DEVICE);`

GPU-GPU communications through MPI



- Memory transfer: GPU \rightarrow GPU \rightarrow CPU $\xrightarrow{\text{MPI}}$ CPU \rightarrow GPU \rightarrow GPU
- GPU aware MPI can help reduce GPU - CPU communication cost
- Communication volume in 3-D is much higher than that in 2-D

Conclusion

- The structured and unstructured algebraic multigrid solvers of HYPRE have been enabled on GPUs by different approaches
- Structured MG solvers have both the setup and solve phases on GPUs, whereas the setup phase of AMG remains on CPUs
- AMG setup on GPUs is work in progress. Major components have been implemented
- HYPRE's abstract memory model enables execution on heterogeneous platforms



CASC

Center for Applied
Scientific Computing

THANK YOU!

Questions & Comments

li50@llnl.gov



**Lawrence Livermore
National Laboratory**

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.